



**State of California
Child Welfare Services/Case Management System**

CWS/CMS County Access to Data (CAD) Architecture

April 5, 2004

Version 2.0





Table of Contents

1.0 System Requirements	1
1.1. Background	1
1.2. Business Requirements	1
1.2.1. General	1
1.2.2. Frequency, Content, Access and Security	2
1.2.3. Workstation Configuration and Installation	3
1.2.4. Training and Ongoing Support	3
1.2.5. Analyst Requirements	4
1.2.6. Technical Requirements	4
2.0 Solution Overview	6
3.0 Solution Components	7
4.0 CAD Database Architecture	9
4.1. CAD Database Environment	9
4.2. CAD Apply Architecture	9
4.2.1. Apply Process Flow Overview	10
5.0 CAD Client Architecture	15
5.1. IBM DB2 CAE	15
5.2. VPN Software	15
5.3. BusinessObjects	15
6.0 CAD Security	16
6.1. Layers and Implementation Methods	16
6.1.1. Network Data Stream	16
6.1.2. RDBMS Security	16
6.1.3. Application Security (BusinessObjects)	17
6.1.4. Connection Types	18
6.2. Data Access Security	19
6.2.1. Key Design Goals and Assumptions	19
6.2.2. Database Schema with Data Access Visibility Flags for County Views	19
6.2.3. DB2 County Views	19
6.2.4. BusinessObjects Security Profiles	19
6.2.5. User Profiles	19
6.2.6. CAD Implementation	20
7.0 CAD Network	21
7.1. County Firewalls and Network Routers	21
8.0 CAD System Utilities	22



8.1.1. Monitoring.....	22
8.1.2. Scheduling.....	22
8.1.3. Backup/Recovery	22
Appendix A – Hardware Configurations.....	23
CAD Test/Development Server.....	23
CAD Production Server.....	24
Appendix B – CAD Server Software	27
Appendix C – CAD Workstation Software.....	28
Appendix D – Listing of Acronyms.....	29
Appendix E – Document Reference.....	31

1.0 System Requirements

1.1. Background

The Child Welfare Services/Case Management System (CWS/CMS) application provides child welfare workers with immediate access to child and family-specific information so caseworkers can make appropriate and timely decisions in child abuse and neglect cases. It also provides child welfare services workers and supervisors with the case management information needed to manage caseloads effectively and efficiently. It is a statewide database that provides the information necessary to monitor, evaluate, and manage California's child welfare services programs, in terms of effectiveness in meeting the needs of the families and children served. It also allows users to meet program goals and mandates.

The primary method of obtaining reports from the CWS/CMS data is by running standard Program Management (PM) reports available in the system. The PM reports are static and are viewed by printed hard copy reports. The secondary method counties use to report on the CWS/CMS data is through the Statistical Analysis System (SAS). SAS is a system that gives users the ability to run ad hoc queries against live production data. However, to run SAS reports, users must be very familiar with the system's complex database structure, which consists of 240+ tables and the associated table fields. In addition, users must understand the SAS programming language in order to write report queries. Well-trained users can write and run SAS reports, but they process at the lowest priority in the system to ensure they do not impact production. Due to these factors, often these reports will take a long time to run and/or must be run during off-shift hours. The factors associated with generating the SAS reports are often much too complicated and burdensome for county personnel.

The formalized request for the County Access to Data (CAD) system was presented in SCR 7390. This SCR detailed the requirements for the system, and *Work Authorization 0003* was approved to authorize the implementation of the solution for CWS/CMS's ad hoc users. The requirements that developed from a Joint Application Requirements session with county and State personnel are detailed below.

1.2. Business Requirements

1.2.1. General

- **Ad hoc Querying Capability:** The solution needed to include ad hoc query, analysis, and reporting that gives users access to data from the CWS/CMS Application Operational Data Store (ODS).
- **Child Welfare Design:** BusinessObjects Universe objects and database table designs needed to be designed with Child Welfare as the priority. This allows CWS/CMS program effectiveness to be measured and allow social workers to make decisions pursuing changes in child welfare based on a program's effectiveness.
- **Electronic Reporting:** The solution needed to be electronic-based, not paper-based.¹

¹ Prior to the implementation of the CAD database environment, all reporting was done using printed report.



1.2.2. Frequency, Content, Access and Security

Frequency:

The CAD ODS refresh rate needed to be, at minimum, weekly.

Content:

The solution needed to:

- Provide a data set containing detailed information
- Support shared report writing
- Provide a single source of CWS/CMS data for statewide and federal reporting requirements
- Include SOC 158 children in placement reports and exclude SOC 158 information unless requested by user
- Provide the ability to integrate other data sources into the CAD environment for future data warehousing needs of Social Services. Currently, no other data sources are used.

Access:

The solution needed to provide users:

- Access to timely ODS data (except for the binary Word documents)
- Statewide visibility of data (for State users)
- Access to data on a 24 hrs/day - 7 days/week basis (24/7)
- Server that is monitored 24/7
- Tools used to access data: BusinessObjects Full Client (Reporter, Explorer and InfoView)

Security:

The solution needed to:

- Allow counties the ability to determine user access to the system within their own county
- Provide a single level of security (user security profile) within a county; i.e., all individuals can view all of the data available to that county
- Support county visibility data constraints imposed in the CWS/CMS mainframe host database county views
- Restrict access to Sealed Cases/Referrals
- Restrict access to Sensitive Cases/Referrals
- Restrict access to Adoptions Information



- Allow additional security requirements that can be reviewed by representatives from county workgroups and State staff and which are managed through the CWS/CMS System Change control process

1.2.3. Workstation Configuration and Installation

Workstation Configuration:

The solution needed to:

- Require minimum or no change to the CWS/CMS Dedicated County workstation image
- Require minimum or no impact on client workstation so that the CWS/CMS Application would not be affected
- Support Windows 95 and Windows 2000

Installation:

The solution needed to:

- Support remote installation
- Contain a Windows 95 and Windows 2000-compatible installation package

1.2.4. Training and Ongoing Support

Training:

The solution needed to:

- Contain an interface requiring minimal tool training
- Provide basic training for reporting tool

Ongoing Support:

The solution needed to:

- Provide ongoing support of tool questions and ad hoc queries
- Keep CAD ODS tables structure in sync with CWS/CMS mainframe host production tables
- Contain online FAQ (Frequently Asked Questions)²
- Provide support available through CWS/CMS Help Desk
- Provide dedicated e-mail for support questions and requests
- Provide updated (timely) user documentation

² Currently, the State of California HHSDC manages and maintains the CAD web site located at: www.hwcws.cahwnet.gov

- Create a User Group to share resources, including face-to-face meetings

1.2.5. Analyst Requirements

The solution needed to:

- Provide end users with the ability to create ad hoc reports
- Provide standard reports with user-defined parameters
- Support historical analysis
- Support static data. Users often refer to as a "snapshot in time"
- Support numerical analysis with counts, summaries, averages
- Provide text and graphical analysis data presentation
- Provide pivot table functionality
- Provide drill-through to detail and drill-up to summary
- Provide user defined time slices - flexibility in report writing and results as regards time periods
- Allow easy manipulation of data sets - allowing for column switching, headings and groupings, slice and dice and filtering
- Share report templates within county
- Provide the ability to save reports and queries locally
- Share report templates among other counties
- Note deltas (changes in data) in reports
- Insert of Legends/Descriptions in report to identify data elements
- Translate SQL query statements into English
- Provide a matrix to map the CWS/CMS Application fields back to the corresponding BusinessObjects Universe objects using business-oriented naming conventions rather than database column name values
- Identify exception to reports for data validation and clean-up
- Provide the ability to query CWS/CMS Application data audit trail information
- Support distribute report results

1.2.6. Technical Requirements

The solution needed to:

- Have minimal or no CPU impact on the CWS/CMS mainframe host environment



- Have minimal network impact between the source mainframe operational data store platform (OS/390) and target CAD database platform (AIX)
- Provide analysts high priority for ad hoc queries in the CAD database system
- Have minimal maintenance of resources and databases required at the county workstation level
- Contain an evolutionary design that embraces new hardware or software components in future functional or performance enhancements
- Focus, where applicable, on commercial off-the-shelf tools rather than custom coding
- Use a centralized design and centrally-located and maintained system to minimize on-site maintenance
- Focus on lower cost alternatives when available and appropriate

2.0 Solution Overview

California's Child Welfare Services (CWS) agencies have extensive reporting requirements to meet county, State, and Federal requirements. CWS agencies currently access CWS/CMS production data through the CAD system using BusinessObjects, a software product that runs on the users' workstations. BusinessObjects is a powerful, full function tool that allows a trained user to create and run ad hoc queries on data stored in CWS/CMS.

The CWS/CMS CAD solution supports several key reporting initiatives for the counties including Division 31, Federal Outcome Measures, and audit issues. Counties use BusinessObjects to query the CAD database in their efforts to manage and track information in areas such as case plans, court dates, medical and dental services, and oversight and administration of staff workload. The ad hoc reporting capability allows county and State users to perform quality assurance functions for the review of case data timeliness, completeness, and compliance with Federal and State policies and standards.

The current CAD user base is 218 users. The State allots each county a number of BusinessObjects licenses, an amount that is based on the relative size of the county and the number of CWS/CMS cases/referrals in that county.

The CAD system is an extension of the CWS/CMS architecture that is a replicated version of the production host data hosted on an AIX (UNIX) platform. The CAD database contains data from the existing CWS/CMS system and is refreshed weekly.

In addition to the statewide view of data, there are logical county-specific database views for all 58 counties in California. With this design, counties are able to query all of their data without affecting the CWS/CMS production environment.




Figure 1.0 System Overview



3.0 Solution Components

CWS/CMS is an online transactional processing (OLTP) system. The CWS/CMS ODS has been optimized for transaction processing and the day-to-day business of child welfare. The ad hoc reporting and analytical processing needs of the CWS agencies require different database optimization criteria that cannot be applied to the CWS/CMS ODS without adversely affecting the performance of the OLTP transactions. Therefore, the CWS/CMS mainframe ODS is not well-suited for ad hoc reporting.

Dynamic report writing requires a flexible database environment capable of processing large volumes of data. It requires indexes specifically designed for the performance improvement of such reports and the data cluster based on county ownership. The CAD system is designed to minimize the impact on the existing production system while providing an optimized database for ad hoc reporting and online analytical processing needs of the CWS agencies.

The CAD environment is built as an extension to the CWS/CMS infrastructure, hosted in the AIX environment, and supported on the existing statewide network. A replicated copy of the production data provides the original data source within the CAD environment. BMC Log Master, in conjunction with custom apply code components, is used to maintain the incremental data updates. Custom database management scripts reorganize the data in the CAD database for improved query performance. Full client BusinessObjects workstations are used to write queries and format the data in the county.

The following diagram (Figure 2.0) illustrates the system components of the logical CAD Architecture.

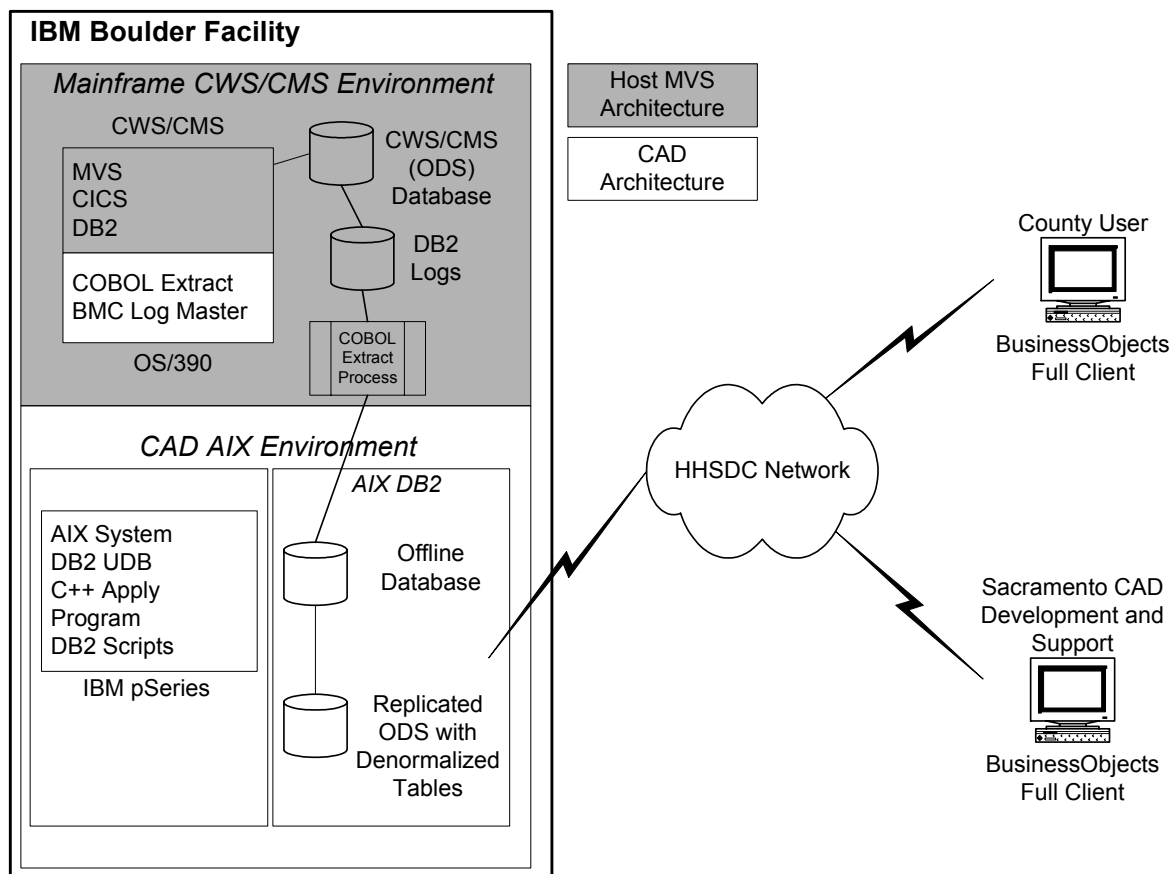


Figure 2.0 - CWS/CMS CAD Logical Database Environment

4.0 CAD Database Architecture

4.1. CAD Database Environment

The CAD solution consists of two databases referred to as the offline CAD Operational Data Store (ODS) and the online CAD ODS. Each of these databases contains one statewide schema and 58 county schemas. The schema “CWSODS” contains a statewide view of the CWS/CMS data. The 58 county schemas “CWSODSnn” (where *nn* is the two digit county identifier) are logical county-based views of the ODS data. The reports and ad hoc queries constructed using BusinessObjects through the CWS/CMS Logical Data Model that is defined in the BusinessObjects universes are transparently executed in the corresponding schema of the online CAD ODS.

The updates to the CWS/CMS mainframe production database are extracted and applied weekly to the offline ODS. Once the weekly data is updated, the offline CAD ODS is catalogued as the online CAD ODS, and the previously catalogued online ODS is re-catalogued as the offline ODS. This toggling is required to allow users access to the online version while the offline version is being loaded, thus providing CWS/CMS users with uninterrupted service.

The databases, schemas, and processes used to apply the CWS/CMS data updates to the offline ODS is referred to as the CAD Apply Architecture.

4.2. CAD Apply Architecture

The CAD apply architecture is designed to provide weekly data updates from the CWS/CMS production database without affecting the BusinessObjects end users. This architecture includes a denormalized version of the CWS/CMS county ownership table that stores the restructured county ownership information. It also includes selected denormalized tables on which the county views are defined to improve performance. The updates from the CWS/CMS production database are directly applied to the statewide schema in the offline CAD ODS. The updates that affect the denormalized tables are extracted using DB2 triggers into a set of temporary tables called the staging area. The denormalized tables are incrementally loaded from the staging area and the statewide schema.

The following diagram (*Figure 3.0*) illustrates the CAD database apply process and its components. Please refer to the subsequent sections for detailed explanations of each of the components.

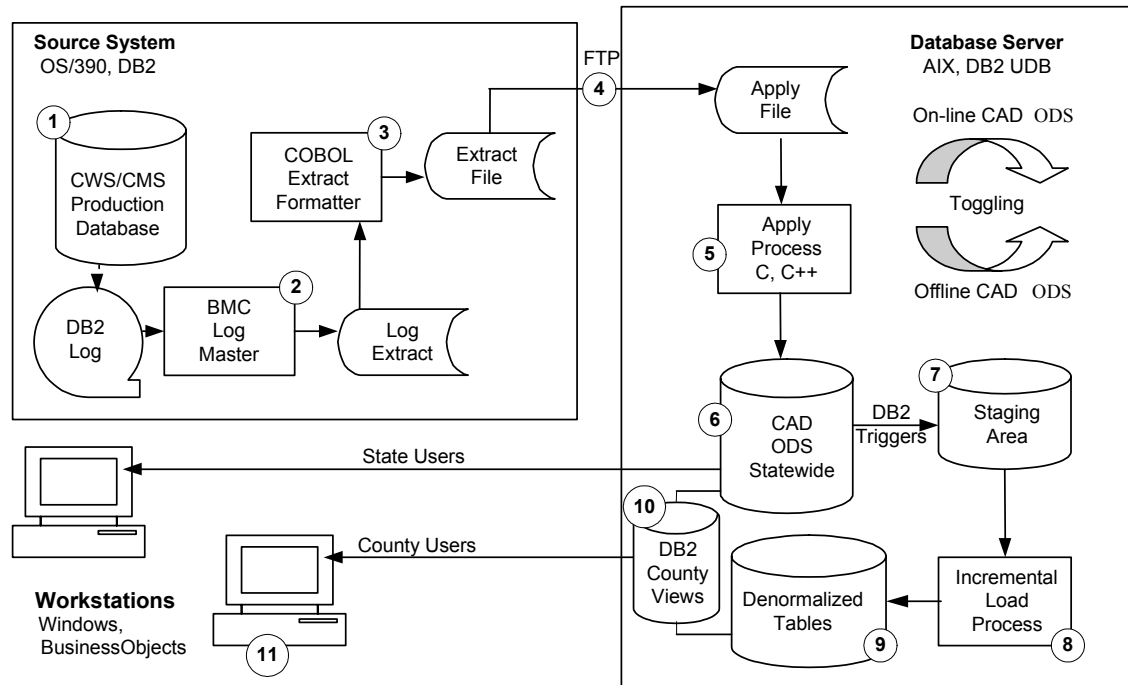


Figure 3.0 CAD Apply Process
See subsequent sections for detailed explanations.

4.2.1. Apply Process Flow Overview

1. CWS/CMS Production Database – a copy of tables from this data serves as the source for the CAD database
2. BMC Log Master - used to capture and extract the updates to the CWS/CMS production tables from DB2 log files
3. COBOL formatter programs - custom-written COBOL programs used to format the extracted apply data files
4. Apply files are transferred to the CAD server via a secure connection using File Transfer Protocol (FTP) on a nightly basis
5. Apply files are processed by custom programs written in C and C++ against the copy of the CAD ODS currently cataloged within the DB2 system catalog as offline
6. Statewide Operational Data Store (ODS)
7. Staging area used to incrementally update the denormalized tables
8. Incremental load process used to update the denormalized tables by processing inserts, updates, and deletes

9. Denormalized tables accessed by users through county views for improved query performance.
10. County DB2 views for optimized query access
11. BusinessObjects Universes for business logic and ease of use

4.2.1.1. *Apply Process (1-5)*

The updates to the CWS/CMS production database (1) are extracted daily from the DB2 log files using the BMC Log Master tool (2) and formatted by a custom COBOL program on the mainframe (3). The COBOL formatter adds the table structure to the extract file created by the BMC Log Master and generates the apply file. The apply file is transmitted via FTP (4) to the CAD server on a nightly schedule.

The apply file is processed by custom programs written in C and C++ (5). These apply programs read the table structure and the extracted records and construct dynamic SQL statements that apply the updates to the offline CAD ODS statewide schema. Because the table structure is embedded in the apply file, the extract/apply process can provide a sophisticated level of generalization. As a result of this generalization, the apply process can accommodate any future changes to the CWS/CMS database without code changes to the extract formatters and apply programs.

4.2.1.2. *Statewide Database Schema (6)*

The statewide schema of the CAD ODS contains a copy of the CWS/CMS production tables made available to county SAS users. The structure of the tables in the statewide schema is the same as the structure of the CWS/CMS production tables with the exception of CASE_T and REFERL_T. CASE_T has two additional columns in the statewide schema, the extract date (column name EXTRACTDATE) and the base 10 Case Number (column name CASE_NO). Similarly, REFERL_T contains two additional columns in the statewide schema, the Extract Date (column name EXTRACTDATE) and the base 10 Referral Number (column name REFERRAL_NO).

The tables in the statewide schema have several additional indexes as compared to the CWS/CMS production tables that are used to support ad hoc queries and provide better query performance. In addition to the indexes, the data in these tables is clustered in the order of the county that owns the record to further improve the ad hoc queries' performance.

4.2.1.3. *County Ownership Structure*

The CWS/CMS production tables replicated in the CAD ODS are classified into four categories depending on the way the county ownership information is maintained. These are:

- Fully-shared tables
- Non-shared tables
- Shared tables with ownership information in the county ownership table
- Shared tables that do not have direct county ownership information

The fully-shared tables are visible to all the counties. The non-shared tables are tables in which each record is owned by one county and is visible only to that county. Each record in the shared tables may be owned by more than one county and is visible to all the counties that own the record. The ownership information for 14 of these shared entities is maintained in the county ownership table in the CWS/CMS production database and subsequently in the CAD ODS. The county ownership information for the remaining shared entities is derived through their parent tables.

The county ownership table in the CWS/CMS production database contains a relational key to one of the 14 shared tables and a type code that identifies the target table. This representation is referred to as a “folded key.” For each folded key, the county ownership table contains 58 attribute flags of the form CTY_nn_FLAG (where *nn* is “01” through “58” for each of the 58 counties). These flags identify the counties that share the target table record identified by the folded key.

The denormalized CAD county ownership table has replaced the CWS/CMS county ownership in the CAD ODS. This table contains the folded key and an attribute that stores the identifier of the county that owns the target table record, identified by the folded key instead of the 58 county flags. If the target table record identified by the folded key is shared by multiple counties -- as indicated by the old county ownership table -- that information will be stored as multiple records, one for each county, in the denormalized county ownership table. In addition, the data in the CAD county ownership table is clustered in the order of the county and the folded key. The data held in the CAD county ownership table is updated through the extract/apply process (see below).

4.2.1.4. Staging Area (7)

The apply process synchronizes the data in the statewide schema of the CAD ODS with the CWS/CMS production data. The denormalized tables used by the county DB2 views are updated through an incremental load process that uses a staging area. The staging area contains one table for each denormalized table and stores the keys of all the records that are inserted, deleted, or updated by the apply process that impact the denormalized tables. During the execution of the apply process in the statewide schema, a set of pre-defined DB2 triggers (actions performed on the databases when certain data conditions are met) extract and insert these keys into the staging area.

4.2.1.5. Incremental Load Process (8)

The incremental load process has two components, the delete component and the insert component. The delete component identifies all the records that are deleted and those that are updated from the staging area and modifies the corresponding records in the denormalized tables. The insert component identifies the records that are inserted and those that are updated from the staging area, derives the county ownership information for these records, and then inserts them in the corresponding denormalized tables.

After the data in the denormalized tables is successfully synchronized with the statewide schema and the CWS/CMS production database, all the data from the staging area is deleted. In addition, the sealed information (such as the sensitive information dealing with adoptions) is deleted from both the statewide schema and the denormalized tables. The offline ODS in which the extract/apply process and the incremental load process are executed is catalogued as the online CAD ODS. The previously catalogued online CAD ODS is re-catalogued as the offline ODS.

4.2.1.6. *Denormalized Tables (9)*

DB2 views determine the visibility of data to a county user. The views derive their county ownership information from the denormalized county ownership table. Even with the restructuring of the county ownership table, the DB2 county views on some of the large and performance-critical tables take a long time to materialize because of the SQL join to the denormalized county ownership table. Denormalized versions of these tables have been created and the corresponding DB2 county views are defined in the CAD database apply architecture. The following are the tables for which denormalized versions have been created:

- ADDRS_T
- CHLD_CLT
- CL_ADDRT
- CLIENT_T
- CLN_RELT
- O_HM_PLT
- PLC_EPST
- PLC_HM_T
- SB_PVDRT

The denormalized versions of the tables above have an additional attribute that identifies the county that owns the record. If multiple counties share a record, then that record will be repeated in these denormalized tables, once for each county. To improve performance, the data in the denormalized tables is clustered by county code. End users access the denormalized tables through DB2 county views.

In addition to the nine tables listed above, several entities have been developed to increase functionality and query performance. These tables were derived from the Data Mart that was originally included in the CAD solution. These tables are updated on a weekly basis through a series of batch scripts that run after the ODS has been refreshed. The following is a list of the additional tables included in the ODS:

- CSDLV90 – This table contains all attempted or completed in-person delivered services performed in the last 90 days.
- REF_MEASR_FACT – This table contains aggregate information about referrals, such as how many minors are listed in the referral, how many cases were opened as a result of the referral, or how many victims are in the referral.
- REF_ABUSE_FACT – This table lists all victims in a given referral by their most serious abuse type as indicated by the hierarchy contained in the SOC 291 report.
- ASGNMT_DIM and ASGNMT_SUP_DIM – These tables contain current or most recent assignment information for cases and referrals.



4.2.1.7. *County Views (10)*

The DB2 county views are based on the statewide schema, the denormalized county ownership table, and the denormalized tables. Each county has its own set of DB2 views that correspond to the CWS/CMS tables in structure. A county's access to data is controlled by these DB2 views and is transparent to the end users. The county and State users are able to use the BusinessObjects universes without having to understand the underlying physical structure of the database.

4.2.1.8. *BusinessObjects Universes (11)*

CAD provides three universes (or semantic layers) for the end user to access the DB2 database. BusinessObjects uses the logic coded in the universe to create SQL that is submitted to the CAD database. Data modeling and relational logic has been incorporated in the universe. The end user does not have to understand SQL or the data model in order to write a business query. The resulting rows are returned to the workstation for formatting and additional processing. The universes reside in the enterprise repository database and are replicated on the workstation. The universes are designed to emulate the CWS/CMS system by organizing the CAD data into logical business units.

The Case Universe represents business elements from the Case folder in CWS/CMS and contains data elements such as placements, court hearings, delivered services, and client demographics. Several additional objects have been provided as well, including date elements such as Federal and State years and quarters.

The second universe, the Referral Universe, represents business elements related to the Referral folder in CWS/CMS and contains business unit data elements such as allegations, placements, and victim demographics. As in Case, several additional elements have been created to provide an easier interface with the data model.

The ODS Universe contains all the data elements available to the CAD user. This universe does not contain any business or relational logic. It is designed to allow an experienced user the flexibility to join the tables in the CAD database to meet their own business needs.

A complete mapping of the Case and Referral Universes and their corresponding CWS/CMS attributes is available to the end users via the State of California HHSDC web site.

5.0 CAD Client Architecture

The CAD client architecture consists of three pieces of software: IBM DB2 CAE, the middleware software used to establish connectivity to the database; the Virtual Private Network software used to encrypt the data stream between the workstation and the server; and BusinessObjects, a full-client end user tool used to build queries and format reports.

5.1. IBM DB2 CAE

IBM DB2 Client Application Enabler is the middleware that allows BusinessObjects to communicate with the CAD DB2 database server. The DB2 database connection information is predefined for the user during the DB2 CAE installation by the Physical Client Application Tool (PCAT) installer.

5.2. VPN Software

CAD data is encrypted for transmission over the wide area network with workstation software and native AIX IPsec utilities on the CAD database server. IRE's Safenet was selected to perform this task for workstations using the Ethernet protocol. Details of this product and its functions are detailed in *Section 6.0, CAD Security*.

5.3. BusinessObjects

The end user query tool, BusinessObjects, was selected for its intuitive interface and flexibility. A requirement for the CAD Project was to provide a system that enables both experienced analysts and first time business users to write ad hoc reports with a degree of simplicity and a minimum of training. BusinessObjects provides a simple interface and "what you see is what you get" formatting.

The CAD solution makes use of custom BusinessObjects universes that were designed and developed by IBM. The universes are organized to provide a user-friendly view of the database. There are three universes in the CAD solution, Case, Referral, and the ODS. The CWS/CMS Application is logically structured into two sections, and the Case and Referral universes mirror these areas. The business logic is built into these universes with the table joins and relationships predefined. The ODS is a straight one-to-one mapping of the CWS/CMS mainframe operational data store and relies on the end user to provide the business logic for the queries.

There is a third method of database query the CAD user can use to write their Structured Query Language (SQL) scripts. This function is called Free-hand SQL in BusinessObjects. Counties are provided with a special DB2 CAE Free-hand connection via PCAT that will only allow them to select data from their views of the database.

The CAD user selects one of these connection methods, universe or Free-hand SQL, and constructs a query in BusinessObjects. The query is submitted by BusinessObjects through IBM DB2 CAE. The CAD database server receives the request, processes the query, and returns the results to the client workstation for report formatting in BusinessObjects.

Data can be exported from BusinessObjects to one of several formats, including Excel, text, and Adobe PDF files. Data and reports can also be saved as HTML files for posting to local county websites.

6.0 CAD Security

6.1. Layers and Implementation Methods

There are a number of layers where security measures are used to prohibit unauthorized access to the CAD system. Within these areas, various combinations of methods provide impermeable security.

Security Layer	Implementation Methods
Network Data Stream	<ul style="list-style-type: none">• IRE's Safenet for encryption between Full Client and AIX for full-client users
RDBMS Security	<ul style="list-style-type: none">• CAE (DB2 Runtime Client) connection (authentication with encrypted password) - per each workstation for full clients• DB2 user authentication
Application Security	<ul style="list-style-type: none">• Assigning security access level for objects through Supervisor administration tool• BusinessObjects Universe Connection (secured connection required for published universes)• BusinessObjects User authentication (logon)
Data Access Security	<ul style="list-style-type: none">• Database Authorization• Database Schema with Flags• DB2 Views• BusinessObjects Security Profiles

6.1.1. Network Data Stream

IRE's SafeNet is used to secure the data stream between the client on the county workstations and the AIX server housing the data warehouse environment databases. Based on the latest IPSec industry standards, the VPN client allows secure Client-to-Client or Client-to-Gateway communication over TCP/IP networks. The security services offered by SafeNet include confidentiality via encryption, packet integrity and authentication via keyed hash, and identity authentication via Digital Signatures and X.509 certificates exchanged during key negotiation.

6.1.2. RDBMS Security

IBM database middleware is installed and configured on each designated CAD user's computer. The Client Configuration Assistant tool allows the definition of data sources, which are alias names used to identify the database a user needs to access. These data source



names are then used to create a BusinessObjects connection. In order to connect and execute queries against the CAD database, the BusinessObjects universes must pass the correct DB2 database connection information to the DBMS including the user ID and password. The DBMS validates or authenticates, and the ID and password authorizes the user to connect to the CAD database. CAD users are only allowed SELECT privileges in the CAD environment. Additional security is enforced through DB2 views and BusinessObjects Supervisor mappings as described in the following sections.

6.1.3. Application Security (BusinessObjects)

Within the BusinessObjects application for the CAD Project, two levels of security are inherently supported and implemented: RDBMS security and BusinessObjects security.

RDBMS security: This allows users to execute BusinessObjects queries through the use of RDBMS account information. This technique is implemented for each of the 59 BusinessObjects Free-hand SQL connections. Each county and the State have a unique Free-hand SQL connection.

BusinessObjects Security: This is provided through the security domain that is set up when a repository is first created. Each repository has its own security domain that contains references to the universes and document domains.

6.1.3.1. Application-Supported RDBMS (DB2) Security

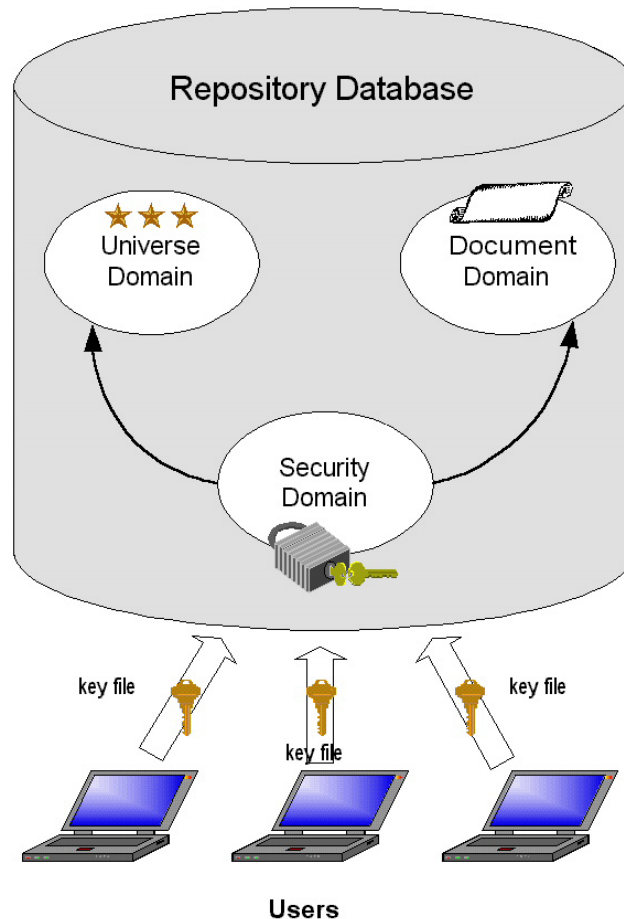
BusinessObjects allows users to connect to a DB2 database using specific data account information for queries and refreshes. CAD Administrators assign user access through a tool called Supervisor. In Supervisor, administrators can choose whether users of a given BusinessObjects universe access the database through the universe's data account, or through their own individual DB2 accounts. In the CAD system, this security technique was applied to the creation of Free-hand SQL connections. Other user connections, (through the BusinessObjects universes) use a combination of DB2 user security (database views) and BusinessObjects table mapping. The mapping is also performed in Supervisor.

6.1.3.2. Application Security Mechanism

BusinessObjects security is provided through the security domain, which is configured when the repository is first created. The repository contains references to the universe and document domains. The security domain also contains information on the identification of BusinessObjects users and the management of the different applications.

The address of the security domain must be recognized by all workstations using BusinessObjects in online mode, so that all users can communicate with the other domains of the repository in a transparent manner. This address is contained in the key file, which is created at the same time as the security domain. This key is distributed to all authorized users.

To access documents, users select a repository to which they have been granted access (via the key file), and enter a valid user name and password. BusinessObjects also allows defining and restricting user access to objects using the administrator tools Supervisor and Designer. CAD uses the Application Security Mechanism as described in this section.



6.1.3.3. Assigning a Security Access Level for Objects

An object can be restricted so that only end users with the appropriate security access level can use it. In this way, certain end users can be prohibited from viewing sensitive or critical information.

Security access levels are assigned to user profiles through Supervisor. The levels are, from highest to lowest: Private, Confidential, Restricted, Controlled, and Public. By default, an object is assigned a Public level.

6.1.4. Connection Types

A connection is a set of parameters that provides access to a database. A connection is made up of three elements:

- The network layer
- The connection name and type
- The database location and account name

BusinessObjects provides three types of connections: secured, shared, and personal. Secured connections are used in CAD for the universe connections. A secured connection is used to

centralize and control access to sensitive or critical data. It is the safest type of connection for protecting access to data. Created with either Designer or Supervisor, secured connections are stored in the security domain of the repository. Any administrator who is granted the appropriate privileges using Designer or Supervisor shares secured connections.

Shared connections are used for Free-hand SQL connections for full-client BusinessObjects users.

6.2. Data Access Security

6.2.1. Key Design Goals and Assumptions

- Prevent unauthorized users from accessing (viewing, printing, exporting, etc.) data from the ODS database, while permitting authorized access to data.
- Make use of the BusinessObjects features that can associate a “where” clause or table alias with a group and automatically apply the “where” clause or alias to queries executed by users who login as a member of a group.
- Data access in the data warehouse environment occurs using SQL and not a procedural language. All security rules must be enforced using SQL supported by an appropriate database design.

6.2.2. Database Schema with Data Access Visibility Flags for County Views

The database schema contains flags that are a materialization of the complex DB2 security views enforced by the CWS/CMS operational environment. The authorization to access county views is granted to the corresponding county’s default user.

6.2.3. DB2 County Views

DB2 views limit data access based on a DB2 user login ID. These views may be materialized at run time (as occurs with the CWS/CMS operational data store) or materialized based on pre-constructed visibility indicators (i.e., database schema with data access visibility flags). Views are materialized only at run time.

6.2.4. BusinessObjects Security Profiles

BusinessObjects provides the opportunity to exclude data based on the BusinessObjects user security profile definition. This is accomplished through SQL “where” clauses or table aliases with restricted “where” clauses.

6.2.5. User Profiles

There is a single user profile defined and supported for the initial release and a single user profile defined in BusinessObjects. The only specification for different data access on a per-user basis is to define access to either a single county or to all counties based on the user’s affiliation.



6.2.6. CAD Implementation

In CAD, a combination of DB2 County Views and User Profiles are employed to ensure a user is only able to access data from their county, unless that user has been defined with statewide data access privileges.

Each county user is assigned to a county group in Supervisor. The county group has separate universe mappings for the Case, Referral, and ODS universes, which relates the group back to a county DB2 view. When a query is written in one of these universes, DB2 alters the SQL to add the county-specific schema to the table names. The query will only return data from a given user's county.

7.0 CAD Network

7.1. County Firewalls and Network Routers

To have connectivity to the CAD system through BusinessObjects, certain protocols and ports of the IP Protocol Suite must be allowed to pass through firewalls and router security devices. Additional network architecture information is provided in the *CWS/CMS Infrastructure Architecture v1.0* document.

The following network ports are required for CAD workstation access to the CAD database: UDP port 500; TCP port 443; IP protocol 50 and 51; and ICMP for pings.

Here is a brief description as to what each protocol will be used for:

- UDP (User Datagram) = IP protocol 17. Used to initiate the VPN secure tunnel negotiation (ISAKMP)
- TCP (Transmission Control) = IP protocol 6. Used to connect to CAD server without encryption (used for troubleshooting)
- SIPP-ESP (SIPP Encapsulated Security Payload) = IP protocol 50. Used to encrypt data (in this case the visible ASCII text of data)
- SIPP-AH (Authentication Header) = IP protocol 51. Used for VPN secure tunnel authentication
- ICMP (Internet Control Message) = IP protocol 1. Used for pinging server (troubleshooting and verifying connectivity)

For all CAD users, VPN profiles are created and stored on the CAD server that serves as the database server and VPN gateway. These profiles are used to authenticate incoming tunnel requests from the CAD BusinessObjects workstations. They include information such as the IP address or the IP subnet that the county CAD workstations belong to.

When a CAD user logs into the BusinessObjects system, the VPN client uses UDP port 500 for ISAKMP key exchange for Phase I of the tunnel negotiation. The CAD server responds on UDP port 500 also to authenticate Phase I.

Phase 2 of the tunnel negotiation and all resulting IP traffic is done on IP protocol 50 and 51. After the tunnel has been established, the CAD BusinessObjects workstation stays authenticated and the tunnel stays up even if the user logs out of BusinessObjects. The tunnel will shut down once the user powers off the workstation. The tunnels will not be initiated unless the user logs into BusinessObjects, pings the CAD server, or tries sending any other type of IP data to the server. The following, generally, describes the overhead for each IP packet for encryption.

Regular IP packet: [IP Header][Data]

Encrypted IP packet: [New IP Header][IPSec Header][IP Header][Data]

The Outer IP header specifies IPSec processing destination. The Inner IP header specifies ultimate packet destination.



8.0 CAD System Utilities

8.1.1. Monitoring

Server availability is monitored through Tivoli Netview. This tool monitors the status of the IP address associated with the CAD server. In addition, Tivoli Distributed Monitoring tools track disk space and capacity, scheduled events, and DB2 availability.

Weekly monitoring of CAD system performance is done by submitting baseline queries via BusinessObjects. The weekly performance of these queries is compared to previous run times for relative performance measures. IBM System Resource Monitoring (SRM) is used to monitor system resources, memory, CPU, and disk space.

8.1.2. Scheduling

CAD weekly processes are scheduled through the AIX native CRON scheduler. All jobs are script-based and outputs from the jobs are evaluated for errors and performance.

8.1.3. Backup/Recovery

The CAD databases are backed up on a weekly basis to a tape library through a scheduled process. An AIX system backup is also taken on a weekly basis. Backup and recovery is managed and controlled using Tivoli Storage Manager (TSM).

In the case of a CAD failure or the need for disaster recovery, the AIX system restoration would be performed and followed by a full restoration of the CAD DB2 databases.



Appendix A – Hardware Configurations

CAD Test/Development Server

Product	Description	Qty
7025-F50	RS/6000 Model F50	1
	CD-ROM Drive	1
	1.44MB 3.5-in Diskette Drive	1
	Base SCSI 6-Pack 1 Kit	1
	Integrated SCSI-2 F/W Adapter 1	1
	Integrated SCSI-2 F/W Adapter 2	1
	Integrated Ethernet Adapter	1
2446	PCI SCSI Adapter to First SCSI 6-Pack Cable	1
2447	16-Bit PCI SCSI SE Adapter to 6 SCSI Bays Cable	2
2493	SCSI-2 F/W PCI RAID Adapter	3
2830	POWER GXT130P Graphics Adapter (PCI)	1
3101	18.2 GB 1" Ultra SCSI Disk Drive Select	1
3104	18.2 GB 1" Ultra SCSI Hot Swap Disk	14
3623	P72 Color Monitor, Stealth Black	1
4106	Select 256 MB (2x128MB) SDRAM DIMM	1
4110	256 MB (2x128MB) SDRAM DIMMs	3
4357	Select 2-Way 604e3 332MHz Processor Card, 2x256KB L2 Cache	1
4959	Token-Ring Adapter	1
5005	Preinstall	1
6206	Ultra SCSI PCI-Bus Adapter	2
6207	Ultra SCSI Differential PCI-Bus	1



Product	Description	Qty
	Adapter	
6519	SCSI Hot Swap 6-Pack	2
8704	Quiet Touch Keyboard, Stealth Black English (UK)	1
8741	3-Button Mouse - Stealth Black	1
9300	Language - English (US)	1
9800	Power Cord - US/Canada (125V, 15A)	1
9996	Internal RAID Indicator	1
7205-311	35 GB Digital Linear Tape Drive	1
1748	Custom Quick Ship Expedite	1
9152	SCSI-2 Cable and Terminator (1m)	1
9200	1 Digital Linear Tape Cartridge Media Kit	1
9300	Language - English (US)	1
9800	Power Cord - US/Canada	1

CAD Production Server

Product	Description	Qty
7038-6M2	Rack Server 1:P SERIES 650	1
2629	4.7GB SCSI2 AUTO-DOCK DVD RAM	1
2848	PWR GXT135P GRAPHICS ACCEL	1
3275	146.8GB 10K RPM U3 SCSI DDRASS	2
4251	BAKPLANE AUTO-DOCKING MEDIA DR	1
4255	INT CAB INTEG CONTR 2SCSI BKPL	1
4651	RACK INDICATOR, RACK 1	1
4959	TOKEN-RING PCI ADAPTER	1



Product	Description	Qty
4962	10/100 MBPS ETHERN PCI ADAP II	1
5005	SOFTWARE PREINSTALL	1
5123	PROC CARD BACKPL 6WAY CONFIG	1
6185	20/40GB INT AUTO-DOCK TAPE DR	1
6203	PCI DUAL CHANNEL U3 SCSI ADAP	1
6287	AC POWER SUPPLY, 1100 W	2
6411	REMOTE I/O LOOP ADAP, PRIMARY	1
6578	U3 SCSI BAKPL HS DISKS	1
8051	2WAY 1.45GHZ PWR4+ PROC CARD	3
8052	4096MB DIMMS EXPRESS CONFIG	3
8058	ENTITLEMENT FC EXPR CONFIG 650M	1
9172	POWER SPECIFY, AC	1
9300	LANG GRP SPEC.-US ENGLISH	1
9911	PWR CORD SPEC.- ALL	1
7311-D10	I/O Drawer 1:I/O DRAWER	1
3147	REMOTE I/O CABLE, 3.5M	2
4651	RACK INDICATOR, RACK #1	1
6006	POWER CONTROL CABLE (SPCN) 3M	2
6230	ADVANCED SERIALRAID PLUS ADAPT	2
6235	32 MBYTE FAST-WRITE CACHE CARD	2
6278	AC POWER SUPPLY, 250W	2
6414	REMOTE I/O LOOP ADAPTER	1
7311	RACK MOUNT I/O DRAW ENCL, 2POS	1
9172	POWER SPECIFY, AC	1
9300	LANG GRP SPEC.-US ENGLISH	1
9911	PWR CORD SPEC.- ALL	1



Product	Description	Qty
3583-L18	Ultrium Scalable Tape Library	1
8002	One Cleaning Cartridge	2
8003	LTO Ultrium LVD Drive Sled	2
8010	20-Data Cartridges	1
9600	Attached to pSeries, RS/6000	1
9704	4.5m VHDCI to HD68 SCSI Cable	1
9800	2.7m Power Cord 125V, 15A - U.S./Canada	1
7133-D40	Advanced SSA Disk Subsystem	4
	(Rack-Mounted)	
987	Rochester Integration	4
8022	50/60Hz AC, 300 VDC Power Supplies	4
8031	Raven Black Drawer Cover	4
8518	One 10K/18.2GB Advanced Disk Drive Module	64
8810	10m Advanced SSA Cable	16
9300	Language - English (US)	4



Appendix B – CAD Server Software

Application Name	Version
Tivoli Management Framework	3.7.1.0
Tivoli Monitoring Client	3.7.0
Tivoli TSM Client	5.1.6.0
Tivoli Storage Manager	5.1.5.0
DB2 UDB EE	7.1.0.77
VisualAge C++ Compiler	5.0.0.0
AIX	5.2.0.0



Appendix C – CAD Workstation Software

Application Name	Version
BusinessObjects Full Client	5.1.3
DB2 CAE	7.1
SafeNet SoftRemoteLT	10.1.1.1

Appendix D – Listing of Acronyms

Acronym	Description
AH	Authentication Header
AIX	Advanced Interactive Executive
DB2	Database 2
CAD	County Access to Data
CAE	Client Application Enabler (Middleware for DB2)
CD-ROM	Compact Disk Read Only Memory
CPU	Central Processing Unit
CRON	Automated Scheduling daemon for AIX
CWS	Child Welfare Services
CWS/CMS	Child Welfare Services Case Management System
DBMS	Database Management System
ESP	Encapsulated Security Payload
FAQ	Frequently Asked Questions
FTP	File Transfer Protocol
HHSDC	Health and Human Services Data Center
HTML	Hypertext Markup Language
IBM	International Business Machines
ICMP	Internet Control Message Protocol
IP	Internet Protocol
IPSec	Internet Protocol Security
ISAKMP	Internet Security Association and Key Management Protocol
ODS	Operational Data Store
OLTP	Online Transactional Processing
PCAT	Physical Client Application Tool
PDF	Portable Document Format
PM	Program Management (in reference to a Report type)
RDBMS	Relational Database Management System
SAS	Statistical Analysis System



Acronym	Description
SIPP	Simple Internet Protocol Plus
SOC	State of California
SQL	Structured Query Language
SRM	System Resource Monitoring
TCP	Transmission Control Protocol
TSM	Tivoli Storage Manager
UDB	Universal Data Base
UDP	User Datagram Protocol
VPN	Virtual Private Network



Appendix E – Document Reference

Document References	Description
URL: http://www.hwcws.cahwnet.gov/county%20logo n/bo_site/CAD_Tools.asp	List of tools available for CWS/CMS CAD BusinessObjects users
<i>BO_Mappings_CaseUniv2_1.xls v2.1</i>	Universe to Host data map for Case
<i>BO_Mappings_RefUniv2_1.xls v2.1</i>	Universe to Host data map for Referral
<i>BusObjMappings221103.xls v2.2</i>	Universe to Host data map for ODS
<i>CWS/CMS Infrastructure Architecture v1.0</i>	Description of overall CWS Infrastructure